

Claims

What is claimed is:

1. A method of managing memory resident queues, said method comprising:

determining whether a memory resident queue is being serviced at a desired level; and

removing one or more messages from the memory resident queue in response to a determination that the memory resident queue is not being serviced at the desired level.

2. The method of claim 1, wherein the removing of the one or more messages causes one or more resources associated with the memory resident queue to be freed.

3. The method of claim 1, further comprising returning at least one message of the one or more removed messages to the memory resident queue.

4. The method of claim 3, wherein the at least one message comprises a set of one or more messages associated with a selected swept identifier.

5. The method of claim 3, wherein said returning is performed in response to a determination that the memory resident queue can accept additional messages.

6. The method of claim 1, wherein the desired level takes into consideration a fullness level of the memory resident queue.

7. The method of claim 1, wherein said determining comprises:

determining whether a size of at least one component of the memory resident queue satisfies a first criterion; and

determining whether a specified portion of messages of a defined type satisfies a second criterion, wherein the memory resident queue is not being serviced at a desired level when the first and second criterion are satisfied.

8. The method of claim 7, wherein the at least one component comprises at least one of a head of the memory resident queue and a tail of the memory resident queue.

9. The method of claim 8, wherein the first criterion comprises being greater than a defined limit of the memory resident queue.

10. The method of claim 7, wherein the specified portion of messages of the defined type comprises one half of a number of get messages.

11. The method of claim 10, wherein the second criterion comprises being less than at least one of a size of a head of the memory resident queue and a size of a tail of the memory resident queue.

12. The method of claim 1, wherein the removing comprises:

allocating storage for the one or more messages to be removed; and

copying the one or more messages from the memory resident queue to the allocated storage.

13. The method of claim 12, further comprising:

identifying the one or more messages being copied with an identifier; and

chaining the identifier to a swept list of the memory resident queue.

14. The method of claim 12, wherein the removing further comprises writing the one or more messages from the allocated storage to at least one storage medium.

15. The method of claim 12, wherein an amount of storage to be allocated is based, at least in part, from where in the memory resident queue the one or more messages are to be removed.

16. The method of claim 15, wherein the one or more messages are to be removed from a tail of the memory resident queue, and wherein the allocating comprises allocating storage for a total number of messages minus a defined number in the tail of the memory resident queue.

17. The method of claim 15, wherein the one or more messages are to be removed from a head of the memory resident queue, and wherein the allocating comprises allocating storage for half a number of get messages being serviced or for a total number of messages minus a defined number in the head of the memory resident queue.

18. The method of claim 1, wherein said determining is performed at one or more time intervals.

19. A system of managing memory resident queues, said system comprising:

means for determining whether a memory resident queue is being serviced at a desired level; and

means for removing one or more messages from the memory resident queue in response to a determination that the memory resident queue is not being serviced at the desired level.

20. The system of claim 19, wherein the removing of the one or more messages causes one or more resources associated with the memory resident queue to be freed.

21. The system of claim 19, further comprising means for returning at least one message of the one or more removed messages to the memory resident queue.

22. The system of claim 21, wherein the at least one message comprises a set of one or more messages associated with a selected swept identifier.

23. The system of claim 21, wherein the returning is performed in response to a determination that the memory resident queue can accept additional messages.

24. The system of claim 19, wherein the desired level takes into consideration a fullness level of the memory resident queue.

25. The system of claim 19, wherein said means for determining comprises:

means for determining whether a size of at least one component of the memory resident queue satisfies a first criterion; and

means for determining whether a specified portion of messages of a defined type satisfies a second criterion, wherein the memory resident queue is not being serviced at a desired level when the first and second criterion are satisfied.

26. The system of claim 25, wherein the at least one component comprises at least one of a head of the memory resident queue and a tail of the memory resident queue.

27. The system of claim 26, wherein the first criterion comprises being greater than a defined limit of the memory resident queue.

28. The system of claim 25, wherein the specified portion of messages of the defined type comprises one half of a number of get messages.

29. The system of claim 28, wherein the second criterion comprises being less than at least one of a size of a head of the memory resident queue and a size of a tail of the memory resident queue.

30. The system of claim 19, wherein the means for removing comprises:

means for allocating storage for the one or more messages to be removed; and

means for copying the one or more messages from the memory resident queue to the allocated storage.

31. The system of claim 30, further comprising:

means for identifying the one or more messages being copied with an identifier; and

means for chaining the identifier to a swept list of the memory resident queue.

32. The system of claim 30, wherein the means for removing further comprises means for writing the one or more messages from the allocated storage to at least one storage medium.

33. The system of claim 30, wherein an amount of storage to be allocated is based, at least in part, from where in the memory resident queue the one or more messages are to be removed.

34. The system of claim 33, wherein the one or more messages are to be removed from a tail of the memory resident queue, and wherein the means for allocating comprises means for allocating storage for a total number of messages minus a defined number in the tail of the memory resident queue.

35. The system of claim 33, wherein the one or more messages are to be removed from a head of the memory resident queue, and wherein the means for allocating comprises means for allocating storage for half a number of get messages being serviced or for a total number of messages minus a defined number in the head of the memory resident queue.

36. The system of claim 19, wherein said means for determining comprises performing the determining at one or more time intervals.

37. A system of managing memory resident queues, said system comprising:

at least one computing unit adapted to determine whether a memory resident queue is being serviced at a desired level; and

at least one computing unit adapted to remove one or more messages from the memory resident queue in response to a determination that the memory resident queue is not being serviced at the desired level.

38. At least one program storage device readable by a machine, tangibly embodying at least one program of instructions executable by the machine to perform a method of managing memory resident queues, said method comprising:

determining whether a memory resident queue is being serviced at a desired level; and

removing one or more messages from the memory resident queue in response to a determination that the memory resident queue is not being serviced at the desired level.

39. The at least one program storage device of claim 38, wherein the removing of the one or more messages causes one or more resources associated with the memory resident queue to be freed.

40. The at least one program storage device of claim 38, wherein the method further comprises returning at least one message of the one or more removed messages to the memory resident queue.

41. The at least one program storage device of claim 40, wherein the at least one message comprises a set of one or more messages associated with a selected swept identifier.

42. The at least one program storage device of claim 40, wherein said returning is performed in response to a determination that the memory resident queue can accept additional messages.

43. The at least one program storage device of claim 38, wherein the desired level takes into consideration a fullness level of the memory resident queue.

44. The at least one program storage device of claim 38, wherein said determining comprises:

determining whether a size of at least one component of the memory resident queue satisfies a first criterion; and

determining whether a specified portion of messages of a defined type satisfies a second criterion, wherein the memory resident queue is not being serviced at a desired level when the first and second criterion are satisfied.

45. The at least one program storage device of claim 44, wherein the at least one component comprises at least one of a head of the memory resident queue and a tail of the memory resident queue.

46. The at least one program storage device of claim 45, wherein the first criterion comprises being greater than a defined limit of the memory resident queue.

47. The at least one program storage device of claim 44, wherein the specified portion of messages of the defined type comprises one half of a number of get messages.

48. The at least one program storage device of claim 47, wherein the second criterion comprises being less than at least one of a size of a head of the memory resident queue and a size of a tail of the memory resident queue.

49. The at least one program storage device of claim 38, wherein the removing comprises:

allocating storage for the one or more messages to be removed; and

copying the one or more messages from the memory resident queue to the allocated storage.

50. The at least one program storage device of claim 49, wherein said method further comprises:

identifying the one or more messages being copied with an identifier; and

chaining the identifier to a swept list of the memory resident queue.

51. The at least one program storage device of claim 49, wherein the removing further comprises writing the one or more messages from the allocated storage to at least one storage medium.

52. The at least one program storage device of claim 49, wherein an amount of storage to be allocated is based, at least in part, from where in the memory resident queue the one or more messages are to be removed.

53. The at least one program storage device of claim 52, wherein the one or more messages are to be removed from a tail of the memory resident queue, and wherein the allocating comprises allocating storage for a total number of messages minus a defined number in the tail of the memory resident queue.

54. The at least one program storage device of claim 52, wherein the one or more messages are to be removed from a head of the memory resident queue, and wherein the allocating comprises allocating storage for half a number of get messages being serviced or for a total number of messages minus a defined number in the head of the memory resident queue.

55. The at least one program storage device of claim 38, wherein said determining is performed at one or more time intervals.

* * * * *